



US007797360B2

(12) **United States Patent**  
**Kretzschmar et al.**

(10) **Patent No.:** **US 7,797,360 B2**  
(45) **Date of Patent:** **Sep. 14, 2010**

(54) **SORTABLE FLOATING POINT NUMBERS**

(75) Inventors: **Klaus Kretzschmar**, Angelbachtal (DE);  
**Nobuyoshi Mori**, Lorsch (DE)

(73) Assignee: **SAP AG**, Walldorf (DE)

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 1195 days.

(21) Appl. No.: **11/400,339**

(22) Filed: **Apr. 6, 2006**

(65) **Prior Publication Data**

US 2007/0240129 A1 Oct. 11, 2007

(51) **Int. Cl.**  
**G06F 7/00** (2006.01)

(52) **U.S. Cl.** ..... **708/204**

(58) **Field of Classification Search** ..... **708/204**  
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,907,435 B2 \* 6/2005 Coulson et al. .... 707/203

7,523,150 B2 \* 4/2009 de Brebisson ..... 708/204  
2006/0179091 A1 \* 8/2006 Carlough et al. .... 708/204

**OTHER PUBLICATIONS**

Cowlshaw, "Densely packed decimal encoding", IEE Proc.-Comput. Digit. Tech., vol. 149, No. 3, May 2002, pp. 102-104.\*  
Draft Standard for Floating-Point Arithmetic P754, IEEE Standards Activities Department (2006).  
Draft IEEE Standard for Floating-Point Arithmetic P754/D0.16.4 Mar. 30, 2006.

\* cited by examiner

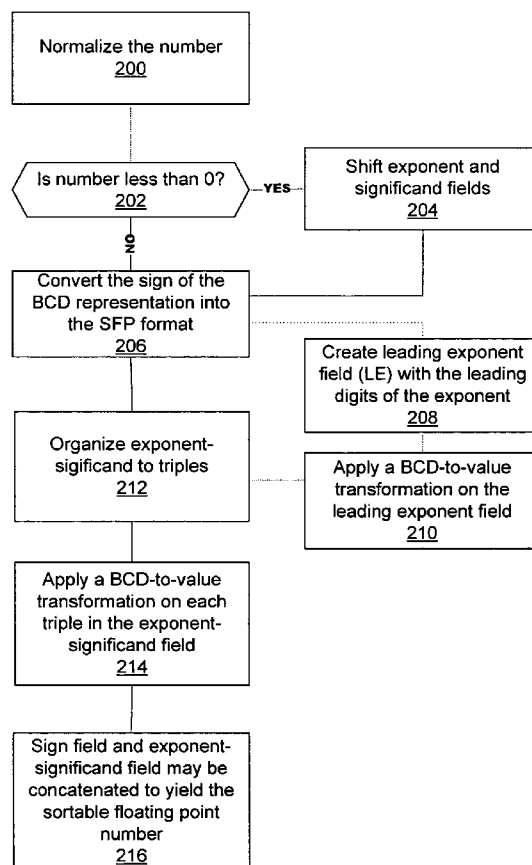
*Primary Examiner*—Tan V Mai

(74) *Attorney, Agent, or Firm*—Kenyon & Kenyon LLP

(57) **ABSTRACT**

The invention comprises methods for manipulating floating point numbers on a microprocessor where the numbers are sortable. That is, the numbers obey lexicographical ordering. Hence, the numbers may be quickly compared using bit-wise comparison functions such as memcmp( ). Conversion may result in a sortable floating point number in the form of a sign, leading bits of the exponent, and sets of digit triples in the form of declets (sets of 10 bits). In a variable-length version, numbers may be compressed by storing the number of trailing zero declets in lieu of storing the zero declets themselves.

**13 Claims, 10 Drawing Sheets**






FIG. 1

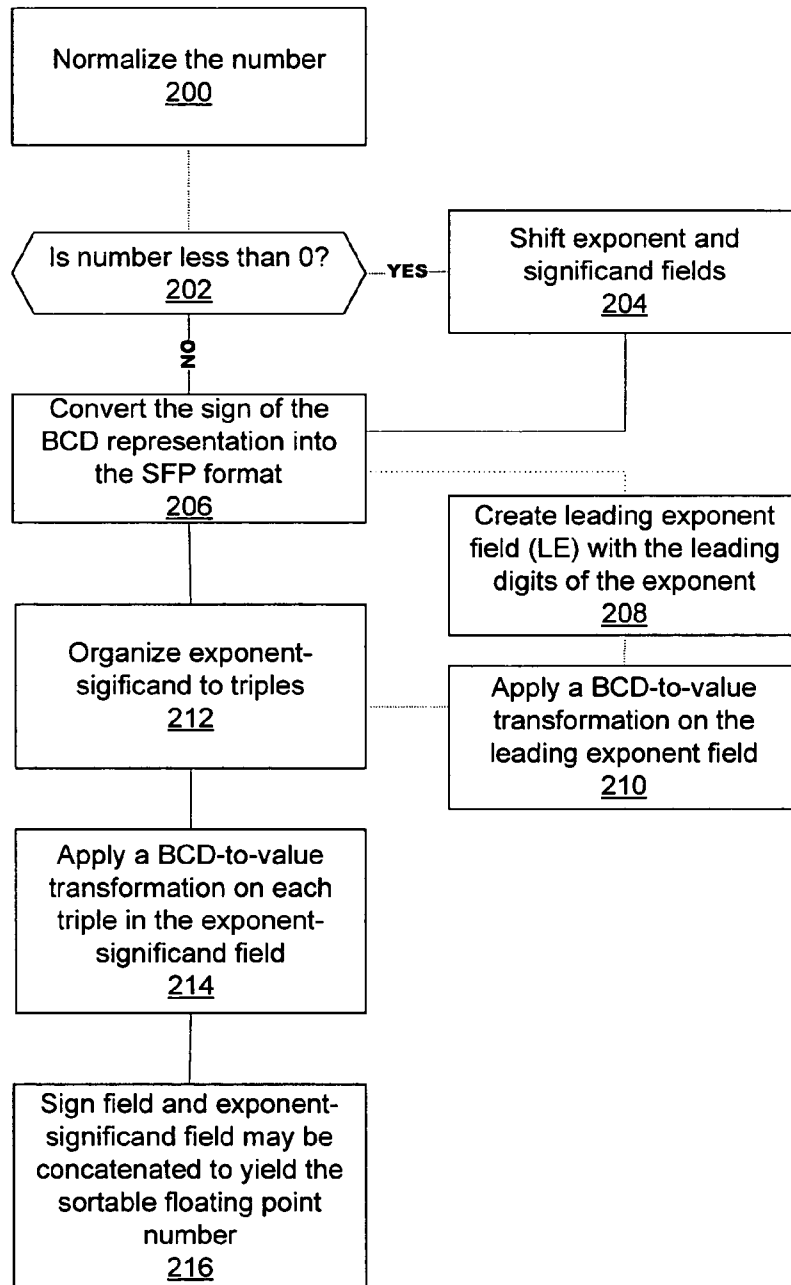


FIG. 2




FIG. 3

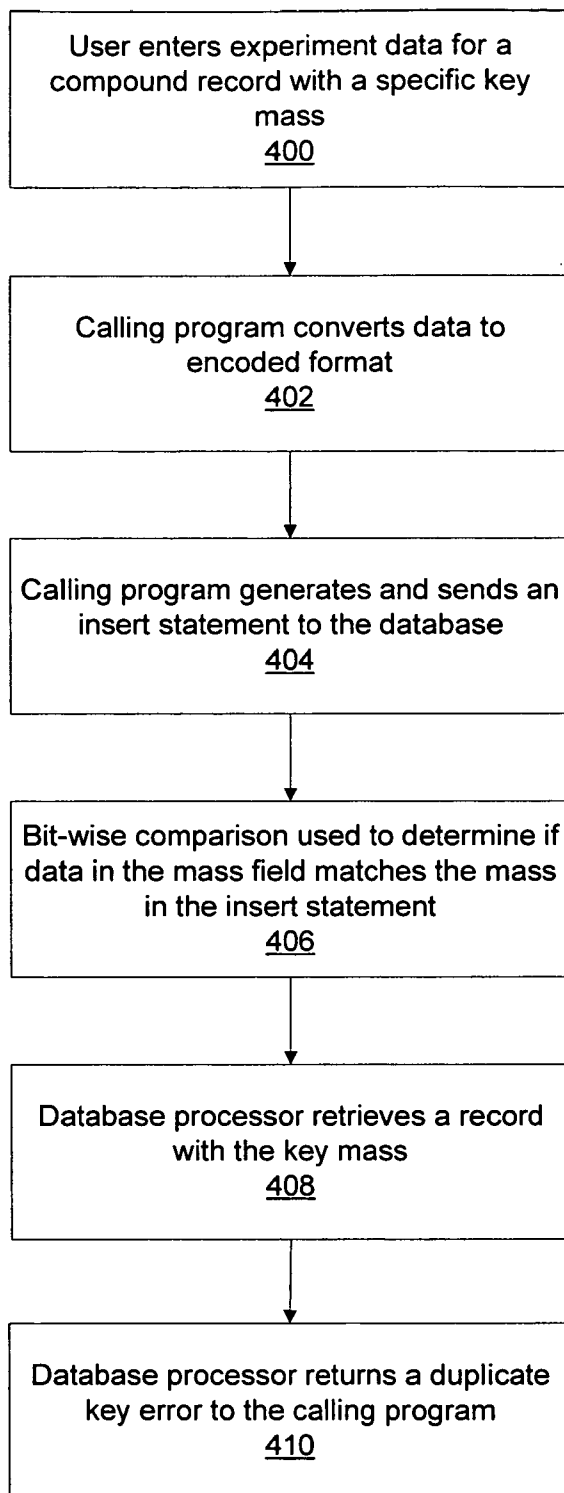


FIG. 4

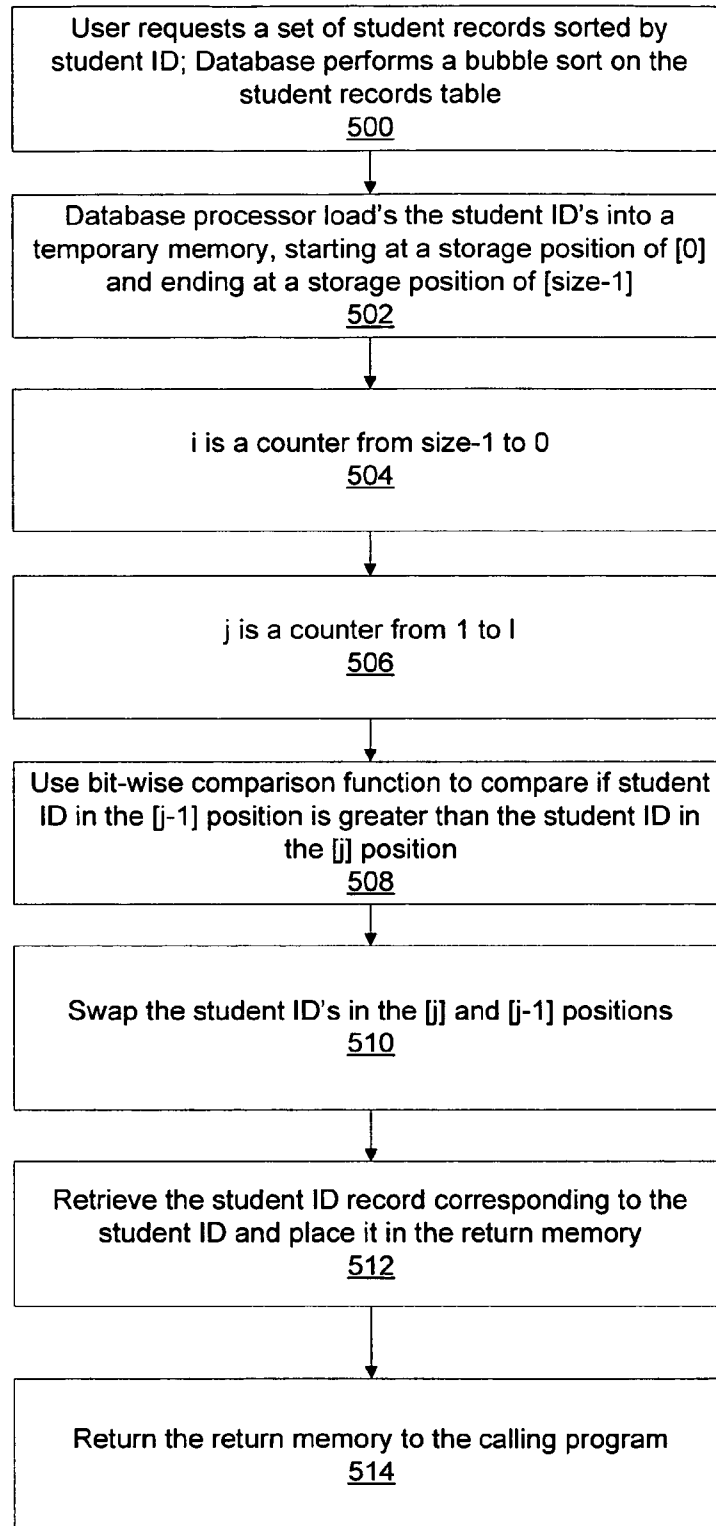


FIG. 5

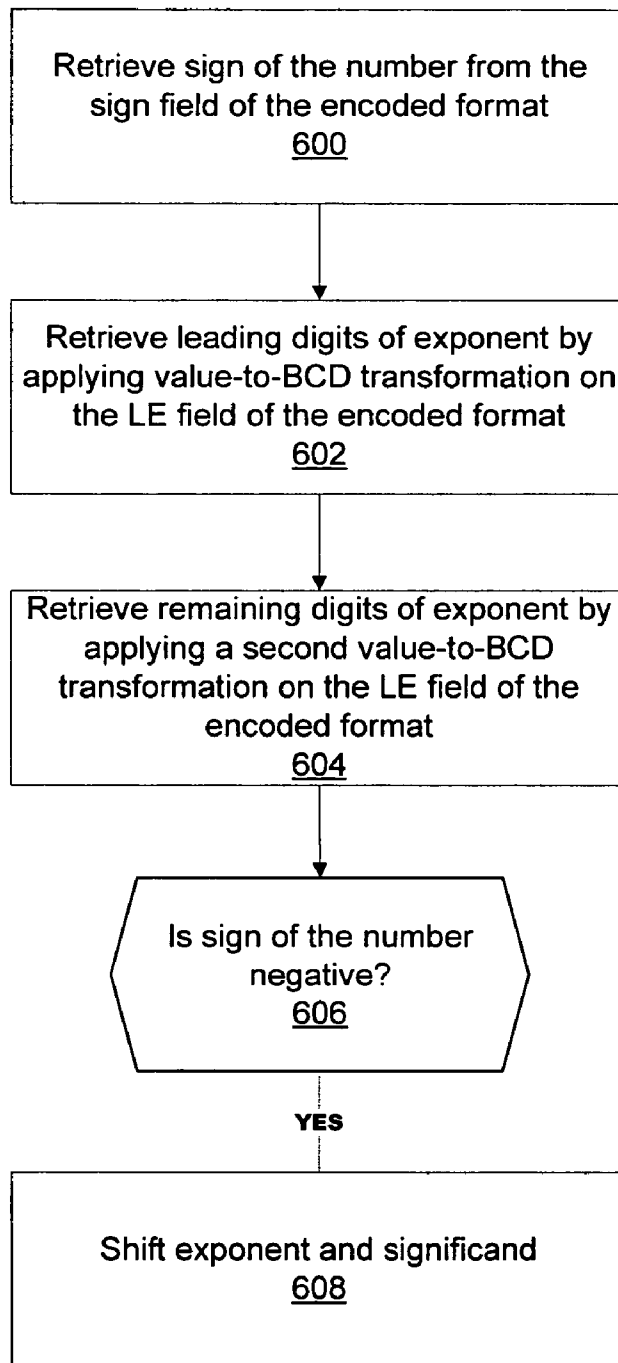


FIG. 6

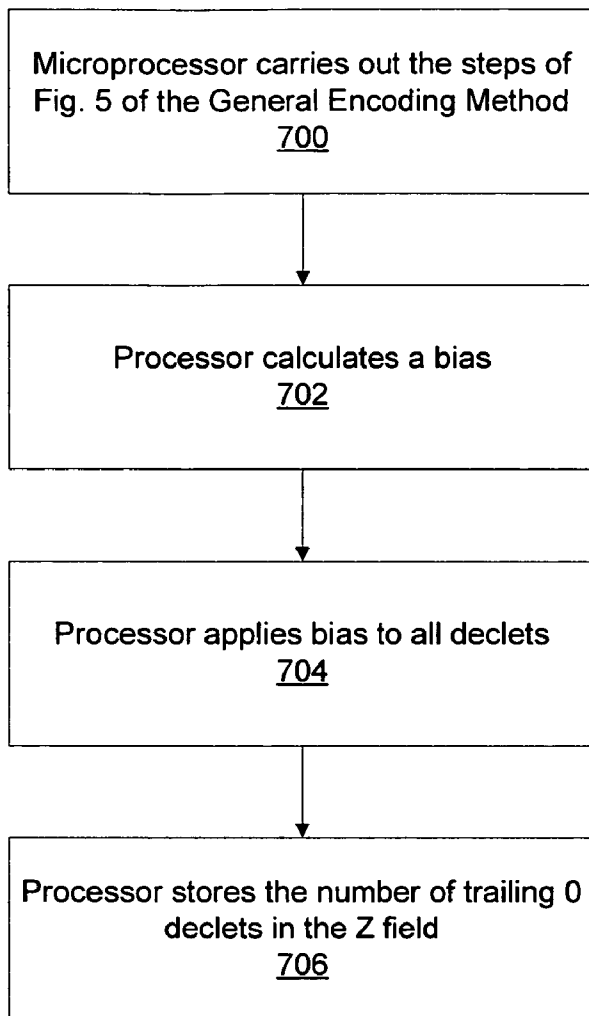


FIG. 7



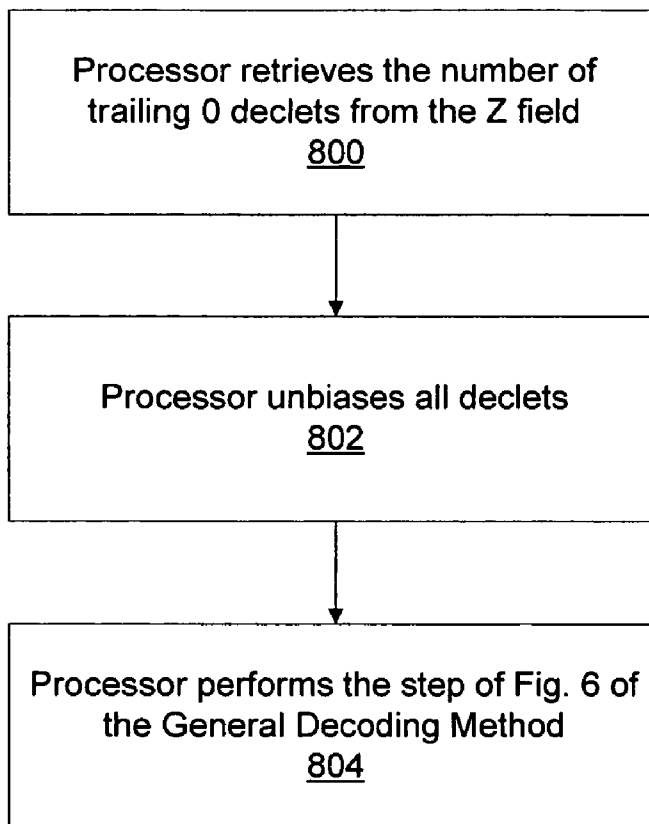


FIG. 8

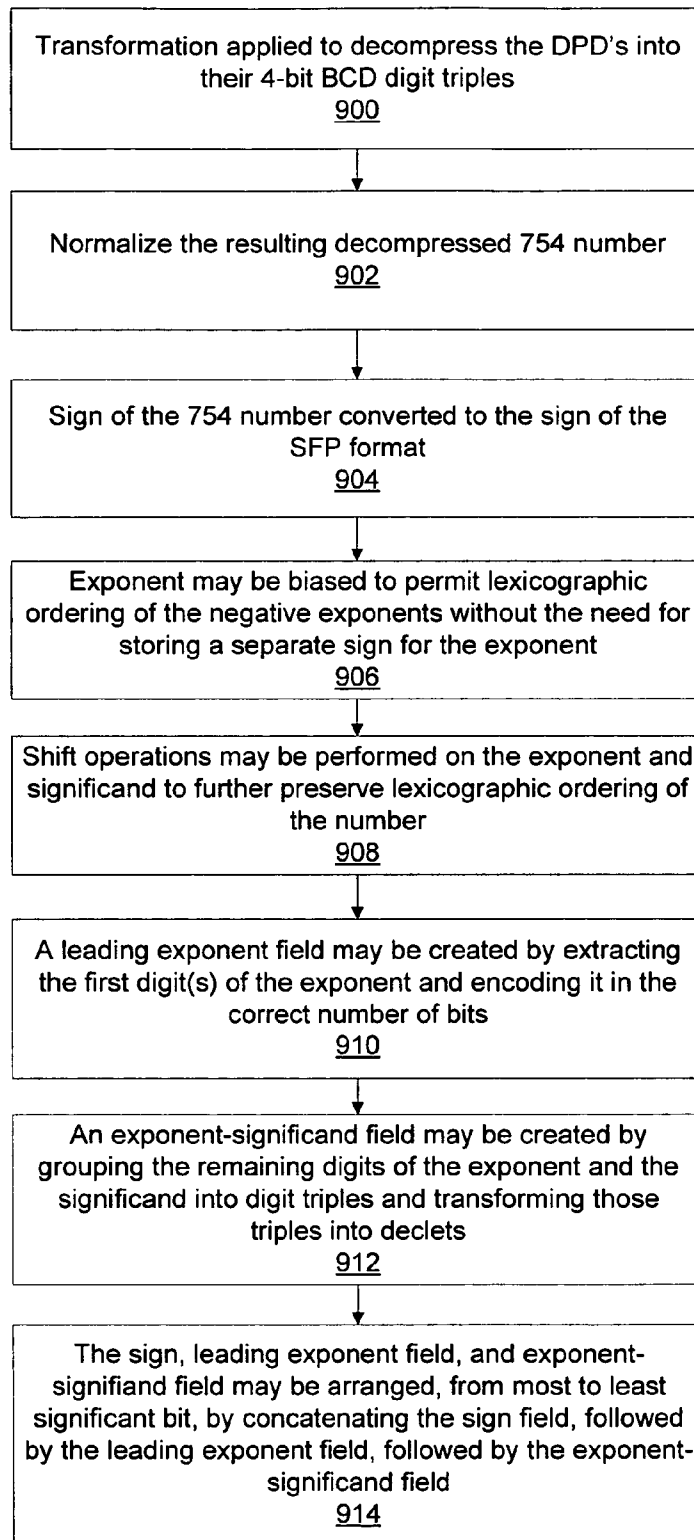


FIG. 9




FIG. 10

## SORTABLE FLOATING POINT NUMBERS

## BACKGROUND

Efficient representation and manipulation of numbers in computer systems presents many challenges for computer scientists and mathematicians. The goal of this art is to find an internal computer representation of numbers that:

Makes efficient use of computer storage (stores the numbers in as few bits as possible.)

Permits easy conversion to numbers for other operations (e.g., ASCII for display to users).

Permits fast numerical computation.

To reconcile these goals, some computers and devices adopt the binary-coded decimal (“BCD”) representation. A BCD representation stores two decimal digits in a binary byte (8 binary bits). Each byte in the BCD representation is broken down into two groups of four bits. One decimal digit may be represented in each group of four bits. Thus, a 10 digit decimal number requires five bytes of storage in the BCD format.

One of the particularly attractive features of the BCD format is the ability to convert BCD numbers to decimal numbers. Since each BCD number represents discrete decimal digits (one digit in four bits), it is relatively simple to determine the value of each group of four binary bits to determine the decimal number of the corresponding decimal digit.

Despite widespread adoption of this representation, microprocessor manufacturers have been slow to formally adopt the BCD representation. BCD manipulations have thus largely been performed in software, as opposed to hardware. What prevents manufacturers from including BCD manipulation hardware in their processors is the lack of a formalized standard for representing floating point numbers in BCD form. Following the codification of a standard, the software manufacturers and database manufacturers can make greater investments in such a format. With widespread adoption comes investment from microprocessor manufacturers in the BCD format.

The emergence of the IEEE 754R representation (“754R”) addresses the standardization of BCD floating point numbers. At the time of this writing, 754R has reached the final stages of standardization. In the near future, it is predicted that virtually all computer manipulations will be performed in the 754R format. Of course, there will be considerable latency between the time that IEEE 754R is ratified and the time that hardware support for this standard is commercially available. Current microprocessors do not possess hardware support for this new standard.

Without dedicated hardware support, 754R presents significant challenges to processor manufacturers and software developers who require specific types of calculations to be fast. For example, it is computationally expensive to compare two 754R values to determine which value has the greater or lesser value. Database queries, which are required to return their responses in an exceptionally short amount of time, cannot afford to absorb the additional processing time needed to manipulate numbers stored in the 754R format. Similarly, data sort operations are extremely expensive. Without sortability in database systems, 754R numbers may also be unable to serve as key fields in tables without implementing expensive software operations to convert the 754R numbers to a sortable format. Thus, there is a need in the art to provide a mechanism for legacy microprocessors and the like to process data in the 754R format but also provide for fast data comparison operations.

## BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 depicts an illustrative system of the present invention.

FIG. 2 depicts illustrative steps of an exemplary embodiment of the present invention.

FIG. 3 depicts illustrative steps of another exemplary embodiment of the present invention.

FIG. 4 depicts illustrative steps of another exemplary embodiment of the present invention.

FIG. 5 depicts illustrative steps of the general encoding method.

FIG. 6 depicts illustrative steps of the general decoding method.

FIG. 7 depicts illustrative steps of the variable-length encoding method.

FIG. 8 depicts illustrative steps of the variable-length decoding method.

FIG. 9 depicts representative steps in a preferred embodiment of the present invention.

FIG. 10 depicts an application accessing a database tier according to an exemplary embodiment of the present invention.

## DETAILED DESCRIPTION

The present invention contemplates systems and methods for efficiently manipulating and comparing floating point numbers to render them amenable to bitwise comparisons. According to an embodiment of the invention, floating point numbers may be converted to a sortable floating point format (“SFP”) in which a first field stores data representing the number’s sign, a second field stores data representing the number’s exponent and a third field represents the number’s significand. Data in these fields may be organized in binary declets, groups of ten bits. Each decllet may represent three decimal digits. Such a conversion permits BCD floating point representations of a certain bit length (e.g., 32-, 64- or 128-bits) to be represented in a SFP format with the same bit length. In another embodiment, numbers may further be compressed by add-biasing all decllets by a predetermined number and storing the number of trailing zero decllets (decllets containing all zeros).

The present invention introduces three exemplary, fixed-length formats with different lengths: decimal32 (32 bit), decimal64 (64 bit), and decimal128 (128 bit) encoding. The invention contemplates fix-length representations of other lengths as well. The present invention also introduces a variable-length encoding which results in a compression factor of up to four for 128 bit numbers.

The present invention provides several advantages over the 754R format. First, software comparisons in the SFP format are markedly improved over those for the 754R format because they permit bit-wise comparison functions, such as memcmp(). Additionally, as noted above, the SFP representation maintains the same bit length as the 754R representation. Thus, a 128 bit number in 754R format can be converted to an SFP format that is 128 bits. Third, SFP numbers may be compressed, thereby significantly reducing storage requirements for such numbers. For example, for the 128-bit format of SFP, a compression factor of four may be realized. Finally, coming on the heels of standardization, the processing system of the present invention, including the possibility to map the 754R numbers to existing native binary database types permits database manufacturers and other developers to use the 754R format without waiting for hardware manufacturers to provide processors with dedicated support for the new format.

FIG. 1 depicts an exemplary system 10 of the present invention. The system 10 contains a database 12 with database table 14 and database processor 16. An interface terminal 18 is shown in communication with the database processor 16. The processor 16 is shown with SFP encoding software 30 and SFP decoding software 32. A user 20 is shown querying the database 12 via database interface program 22 on the terminal 18.

Numbers may be stored in database table 14 in the SFP format. The database processor 16, using SFP encoder 30, may encode numbers arriving to be stored in the SFP format. The database processor 16 may use memcmp( ) to compare numbers within the database, in response to a query. The performance of memcmp( ) is very fast, and thus database queries are likewise very fast.

The following depicts illustrative, general steps carried out by microprocessors encoding and decoding numbers of the present invention.

graphical ordering, a bit value of 0 may be taken to have than a bit value of 1. Therefore, the sign-bit of negative numbers may be 0 and the sign-bit of positive numbers may be 1 to permit lexicographical comparison functions to properly order the negative and positive numbers.

In step 212, the exponent-significand may be organized into triplets, groups of three decimal digits (e.g., 123456 converts to two triples: 123 and 456). In step 214, BCD-to-value transformations may be applied on each triplet in the exponent-significand field (obtains the binary equivalent of the triplet).

A first BCD-to-value transformation may convert the BCD digits into their 10-bit, binary representations. The transform finds the binary equivalent of the value represented by the BCD coding. Using the triplet "012" as an example, the transformation obtains the binary value 00 0000 1100. Table 1 illustrates transformations that may occur in response to input BCD values.

TABLE 1

DECIMAL VALUE OF BCD TRIPLET	BINARY BIT PATTERN OF BCD TRIPLET	VALUE OF BCD BIT PATTERN REPRESENTED IN BINARY	BINARY TRANSFORMATION OF TRIPLET
256	0010 0101 0110	598	10 000 0000
730	0111 0011 0000	1840	10 1101 1010
897	1000 1001 0111	2199	11 1000 0001

To begin, any floating point number may be represented with three fields:

- sign (S)—whether the number is positive or negative,
- exponent—a power of 10, and
- significand—the significant portion of the number.

For example, the number 86000 may be rewritten as  $8.6 \cdot 10^5$ . The sign is positive, the exponent is 5, and the significand is 8.6000.

FIG. 2 is a flowchart illustrating a conversion method for a number coded as sign, exponent and significand, according to an embodiment of the present invention. In step 200, the method optionally normalizes the number to be encoded to align the decimal point of a significand portion and adjust the exponent accordingly. Normalization fixes the radix point of all numbers in a processing system and, therefore, makes a bitwise comparison between values possible. Note that step 200 may be optional. The algorithm may take in a number which has already been normalized.

In step 202, the method determines if the number is less than 0 and, if so, it shifts the exponent and/or significand fields to preserve lexicographical ordering for bitwise comparison (step 204). When two negative numbers have the same exponent, the number with the higher significand actually has the lower value (e.g.,  $-2.2 \cdot 10^1$  is less than  $-1.1 \cdot 10^1$ ) even though the significand 2.2 is greater than 1.1). The shift may overcome this artifact, which can promote bitwise comparisons among negative numbers.

The shift likewise performs an inverse shift operation on the exponent to also permit lexicographical comparison of the negative numbers. For example,  $-2.0 \cdot 10^3$  is greater than  $-2.0 \cdot 10^4$ . Proper lexicographical comparison of these numbers must perform an inverse operation on the exponents so that the encoded format of the exponent 4 is less than the encoded format of the exponent 3.

In step 206, the method may convert the sign-bit of the BCD representation into the SFP format. To promote lexicographical

Note that the binary bit pattern of a BCD triplet, if it were interpreted directly as a binary number, would resolve to an incorrect value (in Table 1, BCD 256 has a value of 598). The BCD-to-binary transformation collapses the BCD representation of a value to its proper value in binary.

Once the exponent-significand field has been created, the sign field and exponent-significand field may be concatenated to yield the sortable floating point number (Step 216).

In alternative embodiments, a sortable floating point number may contain a leading exponent field. In step 208, the method may optionally create a leading exponent field (LE) with the leading digits of the exponent. Each representation of numbers of a particular number of bits contains remainder bits. The invention may group digits, as below described, into groups of 10 bits called decllets. In a 64 bit representation, 60 bits may be taken up by 6 decllets. After reserving one sign bit, three bits may remain to encode a leading digit of the exponent and may encode at most one digit. For 32-bit numbers, one bit may remain (encoding one digit of the exponent) and for 128-bit numbers, 7 bits may remain (encoding three digits).

In step 210, the method may optionally apply a second BCD-to-binary value transformation on the leading exponent digits, which obtains the binary equivalent of the decimal number represented by the BCD encoding. This transformation may be used to retrieve the binary representation of the correct number of bits of the leading digits of the exponents so that the binary bits may later be included in the final SFP number.

In an alternative embodiment, the first two digits of the first decllet of the exponent-significand field may be the trailing two digits of the exponent. Therefore, a 64-bit number, may have an exponent of 799, 7 represented by the 3 bits of the LE field and 99 encoded as the first two digits of the exponent-significand field.

Consider the foregoing operation in the context of an exemplary number originally encoded as a BCD representa-

tion of the 64-bit normalized number 2.0000000000000000 10<sup>401</sup>. This may be encoded as follows:

TABLE 2

Field	WIDTH		
	1BIT S	L BITS L <sub>E</sub>	K * 10 BITS EXPONENT-SIGNIFICAND
Bit string	1	100	0000001100 000000000 000000000 . . .

In step 208, the processor may retrieve the LE digits. Since the number is 64-bit, the exponent may have 3 digits and may have one digit of the LE field in SFP format. Thus, the exponent may be 401. The LE digit, the first digit of the exponent in 64-bit numbers, is 4. Applying step 210, the processor may simply convert one BCD digit bit pattern into its equivalent binary bit pattern. In the case of one BCD digit, the binary equivalent may be retrieved by taking the three least significant bits of the BCD bit pattern. The BCD representation of the digit 4, 0100, may be converted to its binary representation, 100.

Grouping the remainder of the exponent digits and the significand into triplets by applying step 212 may yield

TABLE 3

FIELD	SIGN BIT	LE DIGITS	EXPONENT-SIGNIFICAND DIGITS
digits	1	4	012 000 000 000 . . .

The first triplet of the exponent-significand may contain the trailing two digits of the exponent, 01 as well as the first digit of the significand, 2 to yield 012. The next three digits of the significand are 000. Hence the next triple is 000, and so on.

The method may encode the all BCD triplets in binary. Since the remaining triples are 000, their binary representations are likewise all zeros.

Four use scenarios may highlight the SFP processing systems and methods of the present invention. First, in a database access scenario, a user may request the database to select all products in a product catalog database which cost more than \$1.50 and less than \$2.00. The user 20 of FIG. 1 may implement this scenario from a kiosk of a retail establishment. The user 20 may be represented with a form on the kiosk, the terminal 18, and may enter the appropriate range values on the form. Turning to FIG. 3, in step 300, from a calling program, the user terminal may generate and send a select statement to a database requesting all product records wherein the price field is greater than 1.50 and less than 2.00. For example, the select statement may be “select all from “Products” where “price”>1.50 and “price”<2.00.” In step 302, the database processor may convert the values 1.50 and 2.00 to the encoded format to arrive at the numbers SFP(1.50) and SFP(2.00) respectively. In step 204, for each record in the products table, use a bit-wise compare function (such as memcmp()) to determine whether the value of the price field is greater than SFP(1.50) and less than SFP(2.00). In step 306, if so, copy the record to an output data set. In step 308, return the output data set to the calling program.

For a table with N products, utilizing the present invention may complete the select statement in a time of 2SN+e+d while a time of 2N would be required for the select using traditional 754R software comparison implementations. S may be the factor of savings gained by comparing two SFP

numbers using memcmp as compared to software comparisons of 754R numbers. E may be the time it takes to encode the numbers 1.50 and 2.00 in the SFP format. D may be the decoding time to decode the price field of each record in the output data set. Thus, using an S factor of 0.1 yields a total savings of approximately 80 percent.

In a second scenario, in a table storing experiment data for different compound mixtures keyed by mass of the compounds, demonstrate that an entry with a value of 10 should not create a new record when a record already exists with a key of 10.0. Turning to the flowchart of FIG. 4, in step 400, the user enters the experiment data for a compound record with a key mass of 10. In step 402, a calling program converts the data (including the key of 10) to encoded format (including SFP(10)). In step 404, the calling program generates and sends an insert statement with a key of encoded(10) to the database. In checking for duplicate keys in the table in step 406, for each table record, use a bit-wise comparison function (such as memcmp()) to determine whether the data in the mass field matches the mass of encoded(10) in the insert statement (note that this comparison excludes examination of the scale variable). In step 408, the database processor retrieves a record with a key of SFP(10.0) (note that in ordinary databases this behavior may not occur because the internal representations for 10 and 10.0 may differ). In step 410, the database processor returns a duplicate key error to the calling program.

In a third scenario, depicted in FIG. 5, the system sorts a set of student records in a table by student ID (using a bubble sort). In step 500, from a calling program, the terminal may generate and send a request to a database for a set of student records sorted by student ID. The database may perform a bubble sort (with a processing time of O(N<sup>2</sup>)) on the student records table, where N is the time it takes to compare two ID's. In step 502, load the database processor may load the student ID's into a temporary memory, starting at a storage position of [0] and ending at a storage position of [size-1], where size is the number of student-records in the database. In step 504, for i is a counter from size-1 to 0. In step 506, for j is a counter from 1 to i. In step 508, if the student ID in the [j-1] position is greater than the student ID in the [j] position, using a bit-wise comparison function (such as memcmp()) for the comparison, then swap the student's ID in the [j] and [j-1] positions in step 510. Note that this algorithm performs comparisons on the order of n2. In step 512, after the sort is completed, for each student ID from the [0] position to the [size-1] position of the sorted, temporary memory, retrieve the student ID record corresponding to the student ID and place it into a return memory. In step 514, return the return memory to the calling program.

Using a time savings factor of 0.1, the above sort may save approximately 99 percent ((0.1N)<sup>2</sup>=0.01N<sup>2</sup>).

A fourth scenario highlights the compression of the variable-length aspect of the present invention. A database having the capacity of storing one billion 128 bit numbers would ordinarily consume approximately 15.3 gigabytes of disk space. Each 128-bit number consumes sixteen bytes of disk space. One billion of such numbers consume sixteen billion bytes, or 15.3 GB of disk space. With a compression factor of four using the variable-length encoding, the same database may be stored in as small as 3.8 GB of disk space.

It is important to note that two scenarios highlight the impact of the present invention. In a two-tiered scenario, the introduction of sortable floating point numbers makes it possible to develop high-level applications that perform operations on 754 numbers while permitting databases to store these numbers. In FIG. 10, an application level tier 1000

performs calculations on 754 numbers within an application **1002**. For example, the application may be a sales order and delivery system. The application level tier may communicate with the database tier **1004** by way of a database access subsystem **1006** within application **1002**. Application **1002** may pose a query **1012** to the database for all products costing \$1.50. Within the database **1004**, table **1008** may be accessed in response to a query from application **1002**. Database table **1008** may contain field **1010**.

Field **1010** must utilize data types native to the database **1004** because databases are relatively inflexible with rich data types. Databases typically contain 4-bit, 8-bit, and 16-bit binary data types for storing arbitrary, richer data types. It may be possible to map 754 numbers into these types directly. However, doing so renders databases unable to perform queries on the key fields to retrieve the requested data. The primary comparison function used by databases on the binary data types is memcmp(). Since 754R numbers are not inherently sortable with respect to the memcmp() function, the database would be unable to retrieve data from query **1012** without (1) the database storing the data in a sortable format and (2) query **1012** presenting a sortable number for the database to use in fulfilling its query. Thus, introduction of the sortable floating point numbers permits application developers to begin development of applications using the 754 number format utilizing existing database systems.

In another scenario, a single tiered approach highlights performance gains realized by the present invention. Database manufacturers may wish to support native storage of 754 numbers. This means that database manufacturers may create rich data types which are able to capture the 754 number format. Databases may either compare 754 numbers natively or store an alternate representation of the numbers. In the first case, a software library may be used. However, the performance penalties in this approach make it unrealistic. In the second case, the database may create an alternate format used for comparisons. Since databases naturally use memcmp() for comparisons, it is advantageous to use an alternative format which is sortable. Introduction of sortable floating point numbers facilitates database support for 754 numbers.

FIG. 6 is a flowchart illustrating a decoding method according to an embodiment of the present invention. This method may generally invert the processes performed in FIG. 2. In step **600**, the method may retrieve the sign of the number from the sign field of the number in the SFP format. The sign, as encoded, may be 0 if the SFP number is negative and 1 if positive.

In step **604**, the method may retrieve the BCD digits of the exponent and the significand field by applying a first value-to-BCD transformation on the exponent-significand field of the encoded format (obtains the BCD equivalent of the binary representation of the dectet). In other words, the first value-to-BCD transformation may convert a dectet of 10-bits into 3 groups of four bits in BCD, each 4 bits representing each of the three digits encoded in binary of the dectet.

In step **602**, the method may optionally retrieve the leading digits of the exponent by applying a second value-to-BCD transformation on the LE field bits of the encoded format (obtains the BCD equivalent of the binary representation of the leading exponent digits). Each number of bit representations may contain varying numbers of bits depending on the bit representation. Thus, 1 bit may be retrieved for 32-bit numbers, 3 bits for 64-bit, and 7 bits for 128-bit numbers. The second value-to-BCD takes as input the LE field bits and outputs BCD digits corresponding thereto.

Optionally, in step **606**, if the sign of the number is negative, then in step **608**, the method may optionally shift the

exponent and significand to restore the shift performed in the general encoding method. A shift may have been performed during encoding to permit lexicographical comparison of significands and exponents of negative numbers; the inverse shift operation may allow the number to be properly evaluated. Once the sign, exponent, and significand have been retrieved in the decoding method, they may be used to construct the original 754R number.

An example of the general decoding method may clarify the invention. The SPF number represented in Table 2 may be decoded into the normalized number 2.0000000000000000  $10^{401}$ . Applying step **600**, the sign may be retrieved. Since the sign field contains a single bit set to 1, the normalized number is positive.

Applying step **602**, the next three bits, **100**, may be converted to the first digit of the exponent. Three bits may be retrieved because 64-bit numbers may contain a LE field of three bits. A second value-to-BCD transform, which may be the inverse of the second BCD-to-value transform described in conjunction with FIG. 2, may be applied to the three bits **100** to arrive at the BCD digit 0100 (which has a value of 4).

Applying step **604**, the first dectet of the SPF number, 0000001100 (which has a value of 12) may be converted to a three digit, BCD triple using a first value-to-BCD transformation to arrive at the original BCD digit triple. The first binary dectet bit pattern (which has a value of 12) thus yields a 12 bit binary bit pattern of three 4-bit BCD digits with the values 0, 1, and 2 respectively. As with the exemplary, first BCD-to-value transformation, the exemplary, first value-to-BCD transformation may inherently convert the binary bit pattern in a dectet to a BCD bit pattern. In a more understandable method, the first dectet may be evaluated to arrive at the decimal value 12. A value-to-BCD transformation may be applied to convert the number 12 into the number 18. Grouping the binary bit pattern for the number 18 (00000010010) into 3 groups of 4 bits may yield the three BCD digits of 0, 1, and 2 respectively. The leading exponent digit, 4, may then be combined with the trailing exponent digits (from the first dectet of the exponent-significand field), 01, to arrive at the full exponent 401. The remainder of the digits may thus be the full significand. The sign, exponent, and significand may then be combined in the 754R format.

Other embodiments of the present invention may provide further efficiencies by coding a floating point representation of a number in a variable length format. The variable length coding format can conserve resources within a computer system by storing coded numbers with higher data packing ratios than are permissible with the conventional IEEE 754R coding scheme or with the coding scheme of FIG. 5.

The variable encoding method may omit coding of trailing 0 dectets in a significand and instead, may encode a run length of number of 0 dectets. This thereby saves space ordinarily occupied by dectets of trailing zeros.

In this embodiment, the variable SPF format may contain:

- A sign field,
- A leading exponent field,
- An exponent-significand field, and
- A field (Z) that represents the number of trailing 0 dectets.

FIG. 7 represents illustrative steps for encoding numbers using the variable encoding method carried out by microprocessors of the present invention. In step **700**, the microprocessor carries out the steps of FIG. 5 of the General Encoding Method. This yields an SPF number.







13

Similarly, the unbiased shifting operation of the general decoding method may be the inverse of the bias shifting operation of the general encoding method. Since exponents and significands may be shifting if the sign of the original number is negative, the unbiased operation may restore the number to its original form.

In other embodiments, unused bit patterns of the general format may be used to encode the non-numbers, such as infinity, +/-SNaN, and +/-NaN. Likewise, unused bit patterns may be used to encode representations for the number 0. Refer to the table below for further illustrative examples.

TABLE 9

	DECIMAL32	DECIMAL64	DECIMAL128
NAN	0xFDE00000	0xFD48000000000000	0xFB39800000000000000000000000000000
SNAN	0xFD700000	0xFD2000000000000000	0xFB3700000000000000000000000000000000
INFINITY	0xFD6FFFFFFF	0xFCF000000000000000	0xFB3480000000000000000000000000000000
-INFINITY	0x00AFFFFFFF	0x005000000000000000	0x000500000000000000000000000000000000
-SNAN	0x00A00000	0x002800000000000000	0x000280000000000000000000000000000000
-NAN	0x00000000	0x000000000000000000	0x000000000000000000000000000000000000

In some embodiments, a separate scale variable may be used to track the precision of the numbers. In other embodiments, the scale may be encoded in unused space in the encoded format.

Either big endian, little endian, or other byte-orderings may be used. Also, the algorithm may be used to encode and manipulate numbers of other formats, such as conversion to decimal, binary, BCD, or the like.

The following represents a more detailed description of aspects of the present invention. An example of applying the exemplary BCD-to-value transformation follows.

The first example converts a BCD digit triple into its binary equivalent bit pattern.

$$val(bin(d_i d_{i+1} d_{i+2})) = val(bcd(d_i d_{i+1} d_{i+2}) - (val(bcd(d_i d_{i+1} d_{i+2})) DIV 256) * 156 + (val(bcd(d_i d_{i+1} d_{i+2})) DIV 16) * 6.$$

EXAMPLE

Let the BCD-bit-pattern in memory be  $bcd(d_i d_{i+1} d_{i+2}) = 0 \times 123$  (12-bit long) and hence  $bcd(d_{i+1} d_{i+2}) = 0 \times 23$

Evaluating the BCD-bit patterns as a decimal number may yield:

$$val(bcd(d_i d_{i+1} d_{i+2})) = val(0 \times 123) = 291$$

$$val(bcd(d_{i+1} d_{i+2})) = val(0 \times 23) = 35$$

Computing the right-hand-side of the formula may yield:

$$(val(bcd(d_i d_{i+1} d_{i+2})) DIV 256) * 156 = (291 DIV 256) * 156 = 156$$

(DIV denotes the integer division, i.e. it takes only the integer part)

$$(val(bcd(d_i d_{i+1} d_{i+2})) DIV 16) * 6 = (35 DIV 16) * 6 = 2 * 6 = 12$$

$$val(bin(d_i d_{i+1} d_{i+2})) = 291 - 156 - 12 = 123$$

the binary bit-pattern is  $bin(d_i d_{i+1} d_{i+2}) = 0 \times 7B = 0111 1011$

A second example converts two BCD digits into their binary equivalent bit pattern.

$$val(bin(e_0 e_1)) = val(bcd(e_0 e_1)) - (val(bcd(e_0 e_1)) DIV 16) * 6$$

14

EXAMPLE

Let the BCD-bit-pattern in memory be  $bcd(e_0 e_1) = 0 \times 12$

$$val(bcd(e_0 e_1)) = val(0 \times 12) = 18$$

Computing the right hand side of the formula may yield:

$$(val(bcd(e_0 e_1)) DIV 16) * 6 = (18 DIV 16) * 6 = 1 * 6$$

$$val(bin(e_0 e_1)) = 18 - 6 = 12$$

the binary bit-pattern is  $bin( ) = 0 \times 7B = 0111 1011$

Likewise, an example applying the value-to-BCD transformation may be shown as follows.

A third example converts a binary bit pattern into its equivalent BCD digit triple.

$$val(bcd(d_i d_{i+1} d_{i+2}))_2 = val(bin(d_i d_{i+1} d_{i+2})) + (val(bin(d_i d_{i+1} d_{i+2})) DIV 100) * 156 + (val(bin(d_{i+1} d_{i+2})) DIV 10) * 6$$

EXAMPLE

Let the binary delect bit-pattern in memory be  $bin(d_i d_{i+1} d_{i+2}) = 0 \times 2A5$  (10-bit long)

Evaluating the binary delects as a decimal may yield:

$$val(bin(d_i d_{i+1} d_{i+2})) = 677$$

$$val(bin(d_{i+1} d_{i+2})) = 77$$

Computing the right hand side of the formula may yield:

$$(val(bin(d_i d_{i+1} d_{i+2})) DIV 100) * 156 = (677 DIV 100) * 156 = 6 * 156 = 936$$

$$(val(bin(d_{i+1} d_{i+2})) DIV 10) * 6 = (77 DIV 10) * 6 = 7 * 6 = 42$$

$$val(bcd(d_i d_{i+1} d_{i+2})) = 677 + 936 + 42 = 1655$$

the bit-pattern of 1655 is  $0 \times 677$  which is exactly the bcd bit-pattern of 677.

A fourth example converts a binary bit pattern into its equivalent two BCD digits bit pattern.

$$val(bcd(e_0 e_1)) = val(bin(e_0 e_1)) + (val(bin(e_0 e_1)) DIV 10) * 6$$

EXAMPLE

Let the binary pattern in memory be  $bin(e_0 e_1) = 0 \times 23$

$$val(bin(e_0 e_1)) = 35$$

Computing the right hand side of the formula may yield:

$$(val(bin(e_0 e_1)) DIV 10) * 6 = (35 DIV 10) * 6 = 3 * 6 = 18$$

$$val(bcd(e_0 e_1)) = 35 + 18 = 53$$

the bit-pattern of 53 is 0x35 is exactly the bcd bit-pattern of 35.

The following examples show an illustration of biasing a declct during the variable-length encoding method.

For decimal128->p may equal 34

For decimal64->p may equal 16

For decimal32->p may equal 7

If the input number is positive, the biasing formula may read as follows:

$$val(d_{i+2}, d_{i+1}, d_i) \rightarrow val(d_{i+2}, d_{i+1}, d_i) + pDIV3 + 1$$

EXAMPLE

Let val((d<sub>i+2</sub>, d<sub>i+1</sub>, d<sub>i</sub>)) be a declct with the decimal value of 567 and p=34 then compute

$$val(d_{i+2}, d_{i+1}, d_i) + pDIV3 + 1 = 567 + 34DIV3 + 1 = 567 + 11 + 1 = 579$$

If the input number is negative, the formula may read as follows:

$$val(d_{i+2}, d_{i+1}, d_i) \rightarrow 999 - val(d_{i+2}, d_{i+1}, d_i), \text{ where}$$

EXAMPLE

Let val((d<sub>i+2</sub>, d<sub>i+1</sub>, d<sub>i</sub>))=567, then compute

$$999 - val(d_{i+2}, d_{i+1}, d_i) = 999 - 567 = 432$$

The following table contains example encodings of SFP numbers.

Table of encoded decimal64 numbers (precision=16, e<sub>max</sub>=384):

Number	memory content	Memcmp
m <sub>max</sub> 10 <sup>e<sub>max</sub></sup>	fc f7 e7 f9 fe 7f 9f e7	
2.0 10 <sup>3</sup>	c0 30 00 00 00 00 00	>0
1.9 10 <sup>3</sup>	c0 2f 84 00 00 00 00	>0
1.9 10 <sup>2</sup>	c0 07 84 00 00 00 00	>0
1.8999999999999999 10 <sup>2</sup>	c0 07 83 f9 fe 7f 9f e7	>0
1.8 10 <sup>2</sup>	c0 07 20 00 00 00 00	>0

TABLE 10-continued

Number	memory content	Memcmp
5 1.8000000000000001 10 <sup>0</sup>	bf 57 20 00 00 00 00 01	>0
1.8 10 <sup>0</sup>	bf 57 20 00 00 00 00 00	>0
1.8 10 <sup>-1</sup>	bf 2f 20 00 00 00 00 00	>0
1.7999999999999999 10 <sup>-1</sup>	bf 2f 1f f9 fe 7f 9f e7	>0
1.0 10 <sup>-383</sup>	82 5c 00 00 00 00 00 00	>0
0.9999999999999999 10 <sup>-383</sup>	82 57 e7 f9 fe 7f 9f de	>0
10 m <sub>min</sub> 10 <sup>-e<sub>max</sub> + 1</sup>	80 28 00 00 00 00 00 01	>0
0.0 10 <sup>0</sup>	80 00 00 00 00 00 00 00	>0
-m <sub>min</sub> 10 <sup>-e<sub>max</sub> + 1</sup>	7d 6b e7 f9 fe 7f 9f e7	>0
-1.7 10 <sup>-2</sup>	3e 09 2b f9 fe 7f 9f e7	>0
-1.7 10 <sup>-1</sup>	3d e1 2b f9 fe 7f 9f e7	>0
-1.7000000000000001 10 <sup>-1</sup>	3d e1 2b f9 fe 7f 9f e6	>0
15 -1.7999999999999999 10 <sup>-1</sup>	3d e0 c8 00 00 00 00 00	>0
-1.8 10 <sup>-1</sup>	3d e0 c7 f9 fe 7f 9f e7	>0
-1.8 10 <sup>0</sup>	3d b8 c7 f9 fe 7f 9f e7	>0
-1.7999999999999999 10 <sup>2</sup>	3d 68 c8 00 00 00 00 00	>0
-1.8 10 <sup>2</sup>	3d 68 c7 f9 fe 7f 9f e7	>0
-1.8000000000000001 10 <sup>2</sup>	3d 68 c7 f9 fe 7f 9f e6	>0
-1.9 10 <sup>3</sup>	3d 40 63 f9 fe 7f 9f e7	>0
20 -m <sub>max</sub> 10 <sup>e<sub>max</sub></sup>	00 78 00 00 00 00 00 00	>0

Turning to the optional variable-length encoding, The size of the variable-width format may be given as follows:

25 Size=12+I+(k-val(Z))\*10 bits if D>=0

Size=12+I+(k-val(Z)+1000)\*10 bits if D<0

Hence, to determine the size of the format, we need the number of the trailing zero-declcts. This number can be retrieved by the following algorithm:

30 Start at bit position I+12

Get the declct value declct<sub>i</sub>

1. If D>=0 and val(declct<sub>i</sub>)<=p DIV 3 then val(declct<sub>i</sub>)=val(Z)

2. If D<0 and val(declct<sub>i</sub>)>=1000 then val(declct<sub>i</sub>)=val(Z)

35 3. If p DIV 3<val(declct<sub>i</sub>)<1000 then i->i+1

If i<k-1 go to 2.

Where I represents the length of the leading exponent field and k represents the number of declcts of the corresponding, fixed-size representation.

40 The following represents examples of variable-length encoded SFP numbers.

TABLE 11

Number	memory content (in hex)	memcmp
m <sub>max</sub> 10 <sup>e<sub>max</sub></sup>	FB 34 7F 3F CF F3 FC FF 3F CF F3 FC FF 3F CF F3	
2.0 10 <sup>3</sup>	BD C5 80 B0	>0
1.9 10 <sup>3</sup>	BD C5 F9 00 28	>0
1.8999... 10 <sup>2</sup>	BD C3 78 FF CF F3 FC FF 3F CF F3 FC FF 3F CF F3	>0
1.8 10 <sup>2</sup>	BD C3 72 C0 28	>0
1.80000... 1 10 <sup>0</sup>	BD BE 72 C0 30 0C 03 00 C0 30 0C 03 00 C0 30 0D	>0
1.8 10 <sup>0</sup>	BD DE 72 C0 28	>0
1.8 10 <sup>-1</sup>	BD BB F2 C0 28	>0
1.799999... 10 <sup>-1</sup>	BD BB F2 BF CF F3 FC FF 3F CF F3 FC FF 3F CF F3	>0
1.7 10 <sup>-1</sup>	BD BB EC 80 28	>0
1.70000... 1 10 <sup>-2</sup>	BD B9 6C 80 30 0C 03 00 C0 30 0C 03 00 C0 30 0D	>0
1.7 10 <sup>-2</sup>	BD B9 6C 80 28	>0
m <sub>min</sub> 10 <sup>-e<sub>max</sub> + 1</sup>	80 52 80 C0 30 0C 03 00 C0 30 0C 03 00 C0 30 0D	>0
0.0 10 <sup>0</sup>	80 00 00 C0	>0
-m <sub>min</sub> 10 <sup>-e<sub>max</sub> + 1</sup>	7B 3B FE 7F 9F E7 F9 FE 7F 9F E7 F9 FE 7F 9F E7	>0
-1.7 10 <sup>-2</sup>	3D CC 92 BF C8	>0
-1.7 10 <sup>-1</sup>	3D CA 12 BF C8	>0
-1.70000... 1 10 <sup>-1</sup>	3D CA 12 BF 9F E7 F9 FE 7F 9F E7 F9 FE 7F 9F E6	>0
-1.799999... 10 <sup>-1</sup>	3D CA 0C 80 00 00 00 00 00 00 00 00 00 00 00	>0
-1.8 10 <sup>-1</sup>	3D CA 0C 7F C8	>0
-1.8 10 <sup>0</sup>	3D C7 8C 7F C8	>0
-1.799999... 10 <sup>2</sup>	3D C7 8C 80 00 00 00 00 00 00 00 00 00 00 00	>0
-1.8 10 <sup>2</sup>	3D C2 8C 7F C8	>0

TABLE 11-continued

Number	memory context (in hex)	memcmp
-1.80000 . . . 1 10 <sup>2</sup>	3D C2 8C 7F 9F E7 F9 FE 7F 9F E7 F9 FE 7F 9F E6	>0
-1.9 10 <sup>2</sup>	3D C2 86 3F C8	>0
-m <sub>max</sub> 10 <sup>emax</sup>	00 07 80 00 00 00 00 00 00 00 00 00 00 00 00	>0

As noted, the foregoing embodiments find application for recoding numbers coded according to the IEEE754R format. The basic IEEE 754R format consists of 4 fields:

- 1-bit sign field S.
- 5-bit "combination field" G which contains two leading exponent digits and one leading significand digit.
- 8/10/14-bit "following exponent field" F (for decimal32/decimal64/decimal128 formats) which contains the remaining exponent digits.
- 20/50/110-bit trailing significand field T of densely packed declets (DPD) containing 3 decimal BCD digits. The bcd digits are right-aligned.

exponent to preserve the original value of the number as the raddix point is moved. For example, the number 2000.000000000000 10<sup>-1</sup> may be normalized to yield 2.000000000000000 10<sup>3</sup>. Normalization permits the exponents of the present invention to be properly compared in the encoded format. Although, the normalization step may be optional if numbers arrive to be encoded having already been normalized.

In step 904, the sign of the 754 number may be converted to the sign of the SFP format. If the 754 number is negative, the sign may be 0. Otherwise, the sign may be 1.

TABLE 12

Field	WIDTH			
	1BIT S	5 BITS G	6/8/12 BITS F	20/50/110 BITS trailing significand T
Bitstring	S	G <sub>0</sub> . . . G <sub>4</sub>	F <sub>2</sub> . . . F <sub>w+1</sub>	DPD(d <sub>3-k-3</sub> , d <sub>3-k-2</sub> , d <sub>3-k-1</sub> ) <sup>k-z</sup> <sub>10</sub> . . . DPD(d <sub>2</sub> , d <sub>1</sub> , d <sub>0</sub> ) <sup>1</sup> <sub>10</sub>
most/least significand bit		msb . . . lsb	msb . . . lsb	msb . . . lsb

The three fields S, G, T destroy the byte-wise binary sortability of the IEEE 754R encoding. Because exponent data is distributed across both the G and F fields but can be interrupted by significand values also contained in the G field, a bitwise comparison of two numbers is not possible. Furthermore, the property to store scale information also can prevent a sortable encoding.

Embodiments of the present invention propose to recode IEEE754R into the SFP format described above. To recode this data, exponent values from the G and F fields are reorganized into a continuous exponent field. Similarly, significand data from the G and T fields are reorganized into a continuous significand field. The exponent values are appended to the significand values with the exponent value occupying more significant bit positions than the significand values. Having reorganized the data as such, the data is amenable for transformation as discussed hereinabove.

The following represents a preferred encoding method of the present invention. An example showing the encoding of the 64-bit, 754 number 2000.000000000000 10<sup>-1</sup> helps to clarify this embodiment.

According to the IEEE 754R specification, subpatterns of twelve bits (representing three 4-bit BCD digits) are compressed into subpatterns of ten bits called densely packed declets ("DPD"). In a first step 900, a transformation may be applied to decompress the DPD's into their 4-bit BCD digit triples. Reference may be made to the IEEE 754R specification for examples of the decompression transformation. In step 902, the resulting decompressed 754 number may be normalized according to the description above. The normalization step may fix the raddix point to yield one digit of the significand to the left of the raddix point and the remaining digits to the right. The normalization step may also alter the

In step 906, the exponent may be biased to permit lexicographic ordering of the negative exponents without the need for storing a separate sign for the exponent. An exponent bias value of 398 may be added to the exponent. For example, the exponent of 3 may become 401. In another example, an exponent of -397 may become 1.

In step 908, shift operations may be performed on the exponent and significand to further preserve lexicographic ordering of the number. This shift may be performed if the sign of the 754 number is negative. Reference may be made above for an example of this shift operation. In the current example, since the exponent-biased number 2.000000000000000 10<sup>401</sup> is positive, the shift operation may be skipped.

In step 910, a leading exponent field may be created by extracting the first digit(s) of the exponent and encoding it in the correct number of bits. In this example, the first BCD digit of the exponent is 0100, or 4. According to the preferred embodiment, the leading exponent field of 64-bit numbers may contain three bits. Retrieving the three least significant bits of the BCD bit pattern for 4 may yield 100. Thus, the leading exponent field may be 100.

In step 912, an exponent-significand field may be created by grouping the remaining digits of the exponent and the significand into digit triples and transforming those triples into declets. In the example, the remaining BCD digits of the exponent are 0000 and 0001, or the digits 0 and 1. The significand contains the BCD digits 0010, 0000, 0000, . . . , or the digits 2, 0, 0, . . . Grouping the exponent and significand digits may yield the first BCD digit triple of 0000 0001 0010, or 0 1 2. The next triple may be 0000 0000 0000, or 0 0 0. The remaining triples may likewise have zero values.

19

A transformation may be applied to each digit triple as described above. For example:

Since the first digit of the BCD digit triple is 0, a two digit transformation may be applied.

$$\frac{\text{val}(\text{bin}(e_0e_1)) - \text{val}(\text{bcd}(e_0e_1)) - (\text{val}(\text{bcd}(e_0e_1)) \text{DIV} 16) * 6}{6}$$

EXAMPLE

Let the BCD-bit-pattern in memory be  $\text{bcd}(e_0 e_1) = 0 \times 12$

$$\text{val}(\text{bcd}(e_0e_1)) = \text{val}(0 \times 12) = 18$$

$$(\text{val}(\text{bcd}(e_0e_1)) \text{DIV} 16) * 6 = (18 \text{DIV} 16) * 6 = 1 * 6$$

$$\text{val}(\text{bin}(e_0e_1)) = 18 - 6 = 12$$

the binary bit-pattern of the resulting declat may be 0000001100.

Since the remaining BCD triples are 0's, the resulting binary declats may also be 0.

Finally, in step 914, the sign, leading exponent field, and exponent-significand field may be arranged, from most to least significant bit, by concatenating the sign field, followed by the leading exponent field, followed by the exponent-significand field. This may yield:

TABLE 13

Field	WIDTH		
	1BIT	L BITS	K * 10 BITS
	S	LE	EXPONENT-SIGNIFICAND
Bit string	1	100	0000001100 0000000000 0000000000 . . .

Several embodiments of the present invention are specifically illustrated and described herein. However, it will be appreciated that modifications and variations of the present invention are covered by the above teachings and within the purview of the appended claims without departing from the spirit and intended scope of the invention.

What is claimed is:

1. A method of accessing a database from a first processor, said first processor supporting coding of numbers according to IEEE745R format, said first processor accessing said database through a network, the database being controlled by a second processor that is not in compliance with IEEE745R, the method comprising:

receiving, by the second processor, an input number as a query to a field of the database, wherein exponent values and significand values of the input number are represented in BCD form;

encoding the input number by:

converting, by the second processor, a sign of the input number into a sortable representation to yield a sign field; and

organizing, by the second processor, the exponent values ahead the significand values in significance order, for each consecutive triplet of BCD values, transforming, by the second processor, the triplet into binary valuations of the decimal number represented by the triplet, each transformed triplet yielding a binary declat,

coding, by the second processor, the input number in a sortable floating point form represented by a concatenation of the binary declats and a sign bit,

20

comparing, by the second processor, the encoded input number with numbers stored in the field of the database, retrieving, by the second processor, data stored in the database based on the comparison, and

transmitting, by the second processor, the retrieved data to the first processor through the network.

2. The method of claim 1 further comprising compressing the sortable floating point number by omitting from the number trailing declats having zero value and including a length field to identify the number of declats so omitted.

3. The method of claim 1, further comprising, for any BCD values not included in a triplet, transforming the values into a binary valuation.

4. The method of claim 1, wherein a bit length of a transformed number is no greater than a bit length of the input number.

5. The method of claim 1, further comprising, prior to the transforming, adding a bias to the exponent value to accommodate negative exponent values without requiring an exponent sign.

6. The method of claim 1, further comprising, prior to the transforming, adding a bias to the significand value to accommodate negative significand values.

7. The method of claim 1, further comprising, coding a sign of the input number, wherein a sign value of a positive number is greater than a sign value of a negative number.

8. The method of claim 1 further comprising, prior to the transforming, normalizing the input number to place a radix point of the input number at a known position of the significand.

9. The method of claim 1 wherein the first processor supports a user interface from which a user sends a price request to a produce catalog stored in the database.

10. A method of accessing a database from a first processor, said first processor supporting coding of numbers according to IEEE745R format, said first processor accessing said database through a network, the database being controlled by a second processor that is not in compliance with IEEE745R, the method comprising:

receiving, by the second processor, an input number as a query to a field of the database, the input number being a sortable floating point number;

encoding the input number by:

decompressing, by the second processor, the input number;

normalizing, by the second processor, the input number; if the input number is less than 0, shifting, by the second processor, the exponent and significand fields;

converting, by the second processor, the sign of the input number into a sortable representation to yield a sign field;

creating, by the second processor, a leading exponent field;

organizing, by the second processor, an exponent-significand field into triplets with triplet digits from the remaining exponent bits and the significand of the input number;

applying, by the second processor, a bcd-to-value transformation on the triplets in the exponent-significand field to yield declats; and

assembling, by the second processor, the sign, leading exponent field, and exponent-significand field into the sortable floating point number,

comparing, by the second processor, the encoded input number with numbers stored in the field of the database, retrieving, by the second processor, data stored in the database based on the comparison, and

## 21

transmitting, by the second processor, the retrieved data to the first processor through the network.

11. The method of claim 10, wherein the first processor supports a user interface from which a user sends a price request to a produce catalog stored in the database.

12. A method for accessing a database from a first processor, said first processor supporting coding of numbers according to IEEE745R format, said first processor accessing said database through a network, the database being controlled by a second processor that is not in compliance with IEEE745R, the method comprising:

receiving, by the second processor, an input number as a query to a field of the database, the input number being a sortable floating point number in a variable length format;

encoding the input number by:

decompressing, by the second processor, an input number;

normalizing, by the second processor, the input number;

if the input number is less than 0, shifting, by the second processor, the exponent and significand fields;

converting, by the second processor, the sign of the input number into a sortable representation to yield a sign field;

creating, by the second processor, a leading exponent field;

## 22

organizing, by the second processor, an exponent-significand field into triplets with triplet digits from the remaining exponent bits and the significand of the input number;

applying, by the second processor, a BCD-to-value transformation on the triplets in the exponent-significand field to yield declets;

biasing, by the second processor, the declets with an offset bias, from the second most significant declet to the last declet before the trailing 0 declets;

storing, by the second processor, the number of trailing 0 declets in a Z field; and

assembling, by the second processor, the sign field, leading exponent field, exponent-significand field, and the Z field into the sortable floating point number

comparing, by the second processor, the encoded input number with numbers stored in the database,

retrieving, by the second processor, data stored in the field of the database based on the comparison, and

transmitting, by the second processor, the retrieved data to the first processor through the network.

13. The method of claim 12, wherein the first processor supports a user interface from which a user sends a price request to a produce catalog stored in the database.

\* \* \* \* \*